DataFocus OT · The enterprise ontology platform

Let AI into your business. Keep control.

A business ontology gives AI a map of your business it can understand, a governed way in, and actions that a person approves.

`OT = Ontology` A model of your business that people and AI both work on

Leadership wants AI in production.
Employees already use outside AI tools.
IT and legal won’t hand over the data.

DataFocus OT lets you say yes to AI with confidence.

The problem

What enterprise AI lacks is not the model

Context

The model doesn’t know your business: whether “Customer A” in each system is the same company, how “delivered” is counted, or which version of a rule is in force.

Boundaries

Connect AI straight to your data and nothing controls what it sees or changes. Setting permissions for every AI separately doesn’t scale, and the copies soon disagree.

The ability to act

An AI that only answers questions has limited value. Companies need problems followed through, every step within the rules and a person in control.

A business ontology supplies all three.

OT = Ontology

What a business ontology is

A model of one business area, made of five parts

Part	In DataFocus OT
Objects and relationships	Tables from different systems are organized under the ontology, with how they relate and what each relationship means, shown in the ontology graph
Rules and knowledge	The document center holds policies, rules and metric definitions, linked to the data they govern
Indicators	The indicator center, pushing indicators to the right people on a schedule
Actions	Executions (scripts that calculate, check and write back) and business actions on control panels, such as creating a corrective task
Permissions	Your existing users, roles and resource permissions; AI agents are granted access one ontology at a time

Built-in example

Example: the purchase-approval ontology

Objects and relationships

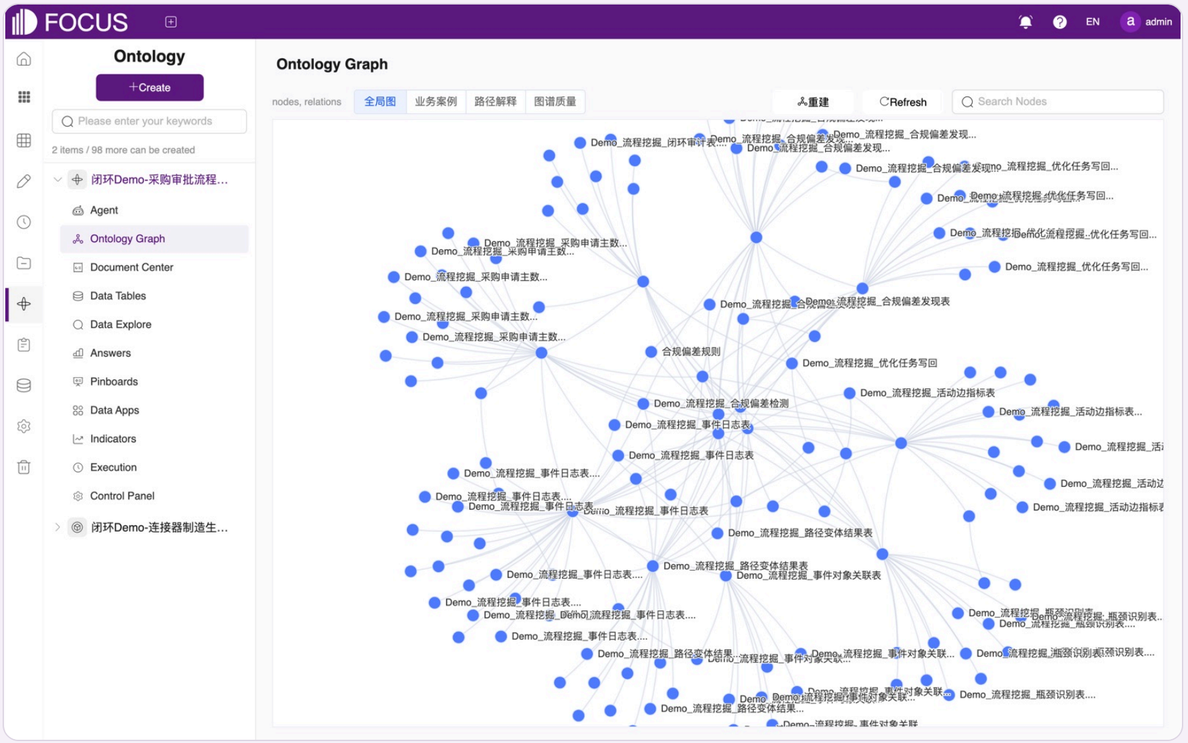
Nine tables, from purchase requests and approval events to compliance findings and audit records, with how they relate

Rules and knowledge

Compliance rules, SLA and bottleneck rules, event-log definitions, and the write-back and audit procedure

Actions and permissions

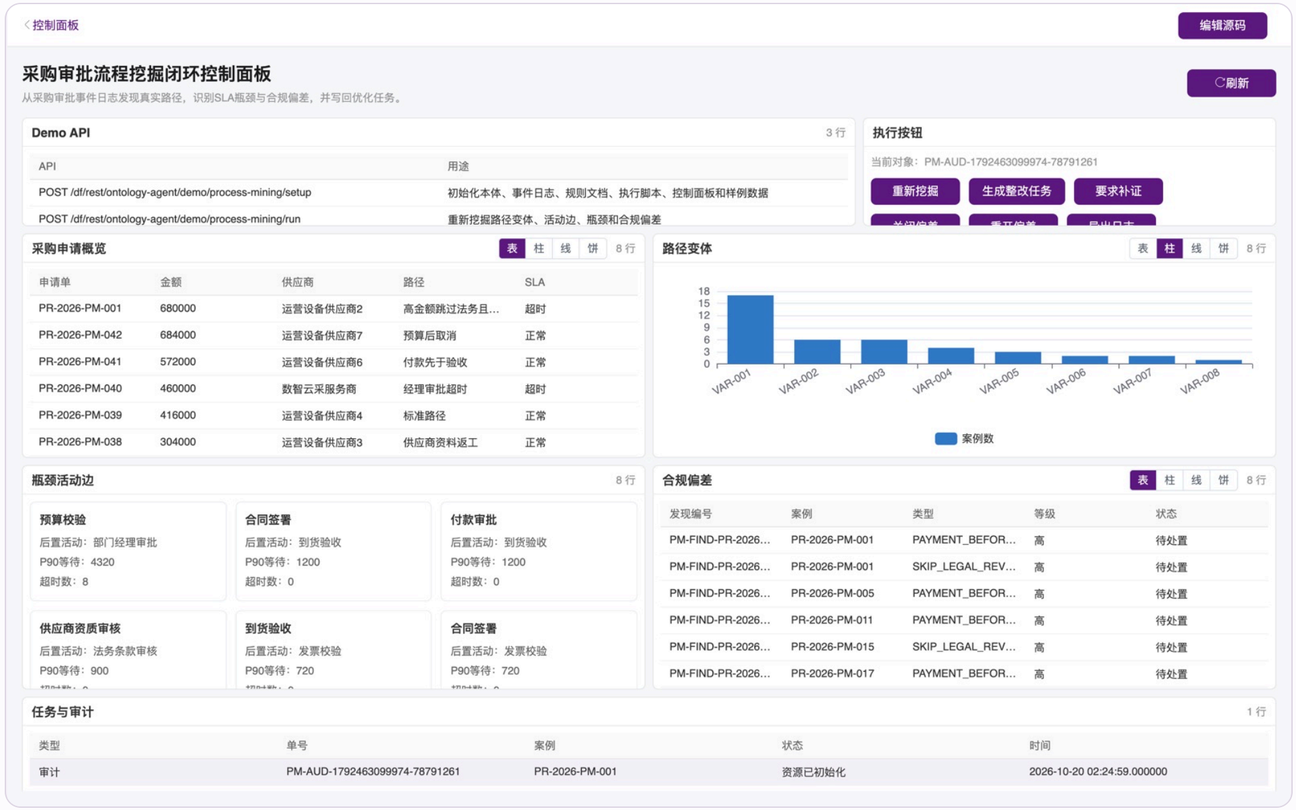
A compliance-check script; create a corrective task, request evidence or close a deviation. Visible only to authorized people and agents



Ontology graph: 147 nodes and 228 relationships

Actions

From finding a problem to fixing it



Everything on one screen

Purchase requests, process variants, bottlenecks and compliance findings, all on the control panel

Act on it directly

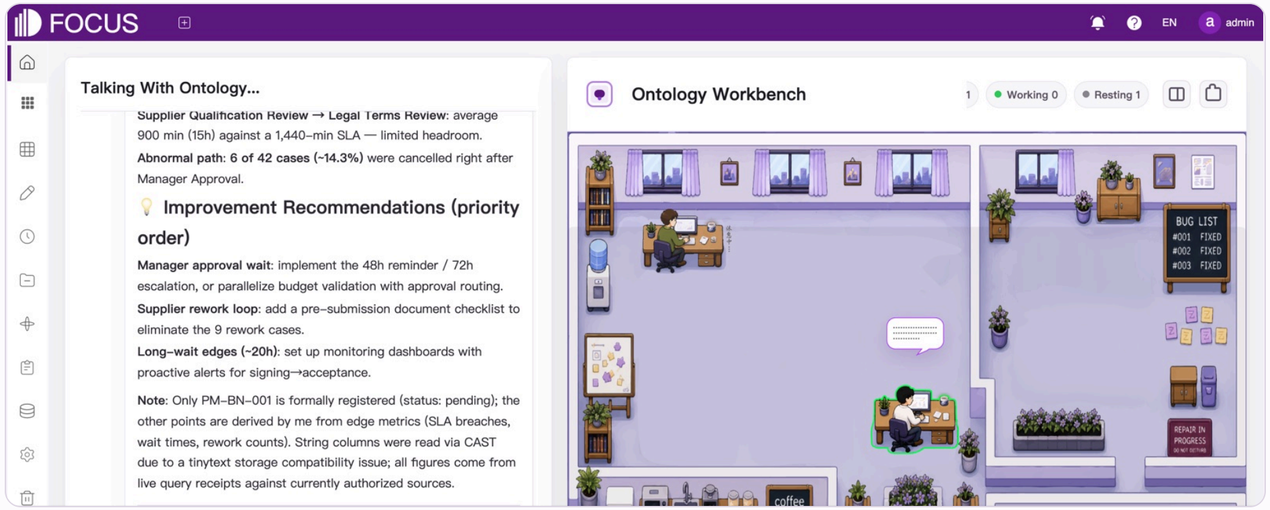
Select an exception to create a corrective task, request evidence or close the deviation

Rules checked automatically

A script finds issues such as large requests that skipped legal review, or payment before acceptance, and writes back the results

Ontology agent

Ask the ontology. Let AI help build it.



Asked about bottlenecks in purchase approvals, the ontology agent analyzes the event log and rules and recommends improvements.

Ontology agent

Ask in plain language. It answers from the ontology's data and documents, with evidence, and can run executions to analyze or write back.

AI helps build the ontology

In the product, AI generates execution scripts and drafts or revises control panels. An external agent can draft an ontology and list questions to confirm; the business owner confirms before publishing.

Agent-drafted ontology: reported by the product team

Evidence you can see

One question, the same figures

Business user · Data Search



AI agent · Claude Code through MCP

The department with the highest total is **Operations Support**, at 7,842,000 yuan.

Operations Support	**7,842,000**
Procurement	**6,966,000**
Digital	**5,786,000**

Basis: the purchase-request master table, with the SQL that ran and a data fingerprint (answer excerpt, translated)

The figures match line for line

Permissions

AI stays in bounds: access follows the ontology

Never beyond the user’s own access

When authorizing an agent, users can tick only ontologies they may access, and nothing is preselected. Ordinary users cannot obtain write access.

Out-of-scope requests reveal nothing

Requests for data outside the grant are refused, and the response does not reveal whether that data exists.

Withdrawal takes effect at once

Once an administrator withdraws access or the user disconnects, the agent’s next call is refused, without waiting for the token to expire.

Every step traceable

Every query returns the SQL that ran, the tables used and a data fingerprint; every agent action is also recorded in an audit log.

Audit log: reported by the product team

Changes

Changes by AI need a person’s approval

When AI wants to create or change something such as a control panel, it can only submit a draft

1
Submit a draft
AI drafts the change

2
Preview on real data
Validated and previewed

3
Create a change plan
The plan cannot be altered

4
Admin approves
Reviewed in the admin console

5
Applied
The approval is on record

A plan that has not been approved cannot be applied, and a forged approval is refused. Submitting the same approval twice does not run it twice.

Value

Value beyond features

Say yes to AI	Turn “not allowed” into “allowed, under control”: access by permission, withdrawable, changes approved. AI gets used, and data stops leaking out the back door.
Stay in control as AI changes	Your business knowledge lives in an ontology on your own servers. You decide which models and agents connect; switch, and you don’t start over.
Value that compounds	The next use case reuses the objects, rules and permissions already there, so each new AI application costs less. One-off pilots are thrown away; the ontology stays.
Accountability when things go wrong	Which data was used, which rule applied, who approved it: all traceable. AI that can be held to account is AI that management can sign off.
Knowledge stays with the company	The rules and practices in experienced employees’ heads become an ontology. People move on; the knowledge stays, and new AI applications start from it.

Why DataFocus

Many claim an ontology: four approaches

Approach	Good at	Falls short
Knowledge graph: entities and relationships extracted from documents	Search and Q&A	Amounts and totals come from the model and are unreliable; no business actions
Metrics or semantic layer built into a data platform	Accurate report figures	No rule documents or business actions; tied to that platform
Data catalog	Describing where data is and what it means	Describes data; plays no part in running the business
Heavyweight ontology platform	Completeness	Expensive and heavy to implement, relying on on-site teams; AI often tied to the vendor’s own products

Why DataFocus

What DataFocus OT does differently

Numbers you can trust

A real compute engine underneath (Doris in Flash; Trino and others in Heavy). Answers come from SQL that actually ran, with a receipt you can check.

Data, rules and actions in one ontology

The data says what happened, the rules say what should happen, and the actions let you deal with it.

Open to AI, governed by default

Any MCP-capable agent connects through standard MCP and OAuth. Permission inheritance, immediate revocation, change approval and query receipts are built in.

Light

Self-hosted, and Flash runs on a single engine. Start with one ontology, without a large on-site team.

Shared by people and AI

Business users’ search analysis, dashboards and scheduled reports run on the same ontology. Invest once, and both sides benefit.

Editions

Two editions

	OT Flash	OT Heavy
Compute and storage	Apache Doris as both in-memory engine and data warehouse	Hadoop, Hive, Trino and Iceberg
Business ontology	Yes	Yes
AI agent access (MCP)	Yes	Yes (reported)
Data pipelines	—	ETL (DataSpring, built on Flink)
AutoML and AI-driven modeling research	—	Yes (reported)
Best for	Teams that want one business area’s ontology running quickly, on a single analytical engine	Enterprises with large data volumes and a big-data platform that need pipelines and modeling

Reported: stated by the product team, not yet verified on a released version

Get started

First step: start with one ontology

1

Pick a business question with data, policy documents and a clear owner

2

Bring the relevant tables and rule documents, and build your first ontology

3

Get business users working on it, then authorize an AI agent to connect

4

Once it works, move on to the next business area

Bring one business question, one table and one policy, and we'll build your first ontology together.

Self-hosted on your own servers · Contact us: datafocus.ai/en/about